

New Methods for Timing Analysis of Transient Events, Applied to Fermi/GBM Magnetar Bursts

Daniela Huppenkothen, Anna L. Watts, Phil Uttley, Alexander J. van der Horst, Michiel van der Klis
Astronomical Institute “Anton Pannekoek”, University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, the Netherlands

Chryssa Kouveliotou

Office of Science and Technology, ZP12, NASA Marshall Space Flight Center, Huntsville, AL 35812, USA

Ersin Göğüş

Sabancı University, Orhanlı-Tuzla, İstanbul 34956, Turkey

Jonathan Granot

The Open University of Israel, 1 University Road, POB 808, Raanana 43537, Israel

Simon Vaughan

X-Ray and Observational Astronomy Group, University of Leicester, Leicester, LE1 7RH, UK

Mark H. Finger

Universities Space Research Association, Huntsville, AL 35805, USA

In order to discern the physical nature of many gamma-ray sources in the sky, we must look not only in spectral and spatial dimensions, but also understand their temporal variability. However, timing analysis of sources with a highly transient nature, such as magnetar bursts, is difficult: standard Fourier techniques developed for long-term variability generally observed, for example, from AGN often do not apply. Here, we present newly developed timing methods applicable to transient events of all kinds, and show their successful application to magnetar bursts observed with Fermi/GBM. Magnetars are a prime subject for timing studies, thanks to the detection of quasi-periodicities in magnetar Giant Flares and their potential to help shed light on the structure of neutron stars. Using state-of-the art statistical techniques, we search for quasi-periodicities (QPOs) in a sample of bursts from Soft Gamma Repeater SGR J0501+4516 observed with Fermi/GBM and provide upper limits for potential QPO detections. Additionally, for the first time, we characterise the broadband variability behaviour of magnetar bursts and highlight how this new information could provide us with another way to probe these mysterious objects.

1. Introduction

Neutron stars present the best test cases for extreme physics in the high-density regime. A long-standing problem in neutron star physics is our lack of understanding of the neutron star interior, in particular, the dense matter equation of state [Lattimer and Prakash 2007]. The detection of quasi-periodic oscillations (QPOs) in the tails of giant flares from Soft Gamma Repeaters (SGRs) has opened up the possibility of studying neutron star interiors using asteroseismology (see Watts 2012 for a review).

SGRs exhibit regular bursts in the hard X-rays and soft γ -rays ($\lesssim 100$ keV), and very rare giant flares with extremely high isotropic equivalent radiated energy of up to 10^{46} erg (see e.g. Palmer et al. 2005). Observations of persistent soft X-ray counterparts showing coherent pulsations with large periods of 5–8 seconds [Kouveliotou et al. 1998, 1999], and the detection of the same periodicities in the tails of the giant flares [Hurley et al. 1999, Palmer et al. 2005], suggested that SGRs are neutron stars. Their behavior is understood within the context of the magnetar model [Thompson and Duncan 1995]: in this paradigm the SGRs are isolated neutron stars with exceptionally strong external dipole magnetic fields, with internal fields that

may be as high as¹ 10^{16} G. Giant flares are powered by a catastrophic reordering of the magnetic field [Woods et al. 2001]. Since this field is coupled to the solid crust, Duncan [1998] suggested that such large-scale reconfiguration might rupture the crust, creating global seismic vibrations that would be visible as periodic modulations of the X-ray and γ -ray flux. This idea was confirmed by the detection of QPOs in the expected range of frequencies ($\sim 10 - 1000$ Hz) in the tails of giant flares from two different magnetars [Israel et al. 2005, Strohmayer and Watts 2005, 2006, Watts and Strohmayer 2006]. If the QPO frequencies can be reliably identified with particular global seismic modes of the neutron star, then they can in principle be used to constrain both the equation of state and the interior magnetic field.

A major obstruction to this field of research is the sparsity of data. Since the launch of the first X-ray and γ -ray instruments, only three giant flares

¹Supported by period and period derivative measurements; see [http : //www.physics.mcgill.ca/~pulsar/magnetar/main.html](http://www.physics.mcgill.ca/~pulsar/magnetar/main.html) for an up-to-date reference list on magnetar spin-down properties

have been observed, with just two having data with a sufficient time resolution to detect QPOs. In trying to overcome this lack of observational constraints, it is therefore a reasonable approach to turn to the much more numerous short SGR bursts with lengths of usually less than a second and luminosities around $10^{40} \text{ erg s}^{-1}$. Hundreds of SGR bursts have now been observed from many magnetars².

To date there has been no systematic search for periodic features in the lightcurves of the SGR bursts. A search for QPOs in a period of enhanced emission with multiple bursts (a ‘burst storm’), from the magnetar SGR J1550-5418, carried out using data from the Fermi Gamma-ray Burst Monitor (GBM), found no significant signals [Kane et al. 2010]. El-Mezeini and Ibrahim [2010] searched a subset of Rossi X-ray Timing Explorer data from SGR 1806-20 for periodic features and found some tentative signals: however, there are several points of concern with regard to their methodology explained in full in the Appendix of Huppenkothen et al. [2013].

Searching for QPOs in transient light curves is a non-trivial task. Standard methods involve Fourier analysis, more specifically the periodogram, defined as the squared amplitude of the Fourier transform of the light curve. The periodogram has several advantageous statistical properties: Poisson noise prevalent in light curves from photon counting experiments results in a flat periodogram with a well-known statistical distribution, making the detection of periodic and quasi-periodic features (as outliers above that distribution) tractable. However, the detectability of QPOs changes significantly in the presence of correlated noise processes and the transient properties of the bursts we are concerned with here. The very nature of a transient event - it has a start, one or more peaks, and an end - complicates the analysis procedure and introduces additional sources of uncertainty, especially in the low-frequency part of the periodogram. For transient events where the shape of the burst envelope is known, many problems arising from the non-stationarity can be solved either analytically [Guidorzi 2011] or via Monte Carlo simulations [Fox et al. 2001]. However, many astrophysical transients such as magnetar bursts do not show a well-behaved burst light curve that is easily reproducible by a simple function. This in itself can be interesting, aside from searching for QPOs: the different burst envelope shapes must be created by a physical process in the source, either in the form of noise processes or non-stochastic emission processes, and characterising the differences may tell

us more about the emission processes at work. The methods and analysis presented in the following summarise more extensive work laid out in Huppenkothen et al. [2013]. We refer the reader to that paper for more details.

2. Variability Analysis

2.1. Monte Carlo Simulations of Light Curves - Advantages and Shortcomings

Monte Carlo simulations of light curves are a standard tool in timing analysis (see for example Fox et al. 2001). The underlying idea is simple: one fits an empirically derived (or physically motivated) function to the burst light curve. One then generates a large number of realizations of that burst profile, including appropriate sources of noise, such as Poisson photon counting noise. The periodograms computed from these fake light curves form a basis against which to compare the periodogram of the real data. For each frequency bin, a distribution of powers is produced. Comparing the observed power in each bin with the distribution of simulated powers in the same bin allows us to make a statement about the probability of the observed power in a particular bin being due to a noise process: if the observed power in a particular bin is a high outlier compared with the distribution of simulated powers in that bin, then the probability of observing the data under the (null) hypothesis of a noise process is $1/N$, where N is the number of simulations performed. If N is large, the observed outlier is unlikely to be produced by the noise process alone.

The Monte Carlo method outlined above is versatile and powerful, but it has limitations. The most important limitation comes from our lack of knowledge of the underpinning physical processes producing the observed light curve. Only if the null hypothesis accurately reflects the data - apart from the (quasi-)periodic signal for which we would like to test - is the test meaningful. If important effects that distort either shape or distribution of the powers are missed, then the predictions made will not be accurate, leading to either spurious detections or real signals not being found.

More often than not, especially in the case of short magnetar bursts, we do not have complete information about the emission mechanism. Short magnetar bursts are extremely diverse, varying in light curve shape as well as burst intensity and duration (see, for example, Göğüş et al. 1999 and Göğüş et al. 2000). There is a fundamental degeneracy in the problem: which features do we consider to be part of the burst envelope, and which a (potential) feature which we do not include in the light curve template? Stochastic processes correlated on different time scales can mimic a quasi-periodic process to the human eye, but

²see e.g. Woods and Thompson 2006, Mereghetti 2008 for overviews or <http://f64.nsstc.nasa.gov/gbm/science/magnetars/> for a collection of SGR bursts observed with Fermi/GBM

clearly, this is not what we wish to detect. Not taking the presence of these features into account can lead to a large number of false positive detections. On the other hand, overcorrecting for features may lead us to detect no features at all. Without detailed knowledge of the burst emission processes, it is impossible to build a reliable model for individual burst light curves to test against when searching for QPOs. We therefore advocate a different method, based on an empirical model of the periodogram.

2.2. Modeling the Periodogram

We give a very short overview of the general principles our method employs. Details can be found in Huppenkothen et al. [2013] and Vaughan [2010]. The Bayesian procedure we employ in modeling the periodogram has three parts: (a) find the preferred broadband noise model to represent the low-frequency part of the periodogram, (b) search the periodogram for the highest outlier and compare this outlier to those distributed by pure broadband noise to find narrow features, (c) search for QPOs in the data, using binned data as well as an identical approach for the model selection in the first step. A step-by-step description can be found in the Appendix of Huppenkothen et al. [2013].

Every step in the analysis follows the same logic: assume a null hypothesis and an alternative hypothesis, compute statistics to summarise the data-model fits for the two different models, generate a sample from this null hypothesis using a Markov Chain Monte Carlo (MCMC) approach, more specifically the MCMC code *emcee* [Foreman-Mackey et al. 2012]. One can then compare the distribution of the relevant statistic derived from the sample generated from the null hypothesis to the observed value of that statistic. If the observed value lies in the high-end tail of the distribution, then it is an outlier with respect to the null hypothesis.

Since the entire procedure rests on the correct choice of broadband model, this is the first step of the analysis. The data are fitted with two continuum noise models, which, by definition of the likelihood ratio test, are required to be nested. The likelihood ratio is the statistic we use to decide which model is preferred by the data. We simulate a large number of fake periodograms from parameter sets drawn from the posterior distribution of parameters, as approximated by a large number of MCMC simulations. Then these fake periodograms are fit with both models again to build a distribution of likelihood ratios from the simple model. We can compute the tail-area probability (p-value) of the observed likelihood ratio to be typical of the distribution (equivalent to asking whether the observed data is sufficiently described by the simpler model) by integrating over the tail of the

distribution. If this probability is lower than a chosen significance threshold, then the data is more likely to be drawn from the more complex model hypothesis, which should then be adopted for the rest of the analysis.

We extensively tested our method on synthetic data generated with known parameters. These simulations confirm that in the limit of white noise, our method matches the predictions from standard Fourier analysis. Furthermore, for more complicated light curves, involving red noise and a burst envelope, our method is conservative in nature at low frequencies, where burst envelope and red noise dominate, but approaches the white noise predictions for high frequencies. The full description of these simulations is given in Huppenkothen et al. [2013].

3. Observations

To test our methodology, we applied it to a small sample of bursts from SGR J0501+4516. Fermi/*GBM* triggered 26 times on this source between 2008 August 22 and 2008 September 03, observing 29 bursts. Two of these (burst IDs *080824054* and *080825200*) had saturated parts, and were therefore excluded from the analysis due to the rather complicated effects saturation can have on periodograms. Following Lin et al. [2011], we used only NaI detectors with a source angle smaller than 50 deg for each of the 24 triggered and 3 untriggered bursts. We use high-time resolution time-tagged event (TTE) data, for which time and energy for each individual photon is recorded. The data were barycentered and channels converted to the mid-energy of each energy bin. The observations were then energy-selected to include only counts between 8 and 100 keV. The lower limit to the energy is set by the detector response [Meegan et al. 2009], the upper limit was found by inspecting energy-resolved light curves and finding no source counts above 100 keV (the counts were consistent with the Poisson distribution expected from counting noise). Burst start times and lengths (T90 durations) were taken from Lin et al. [2011], and are summarized in Table 1 of that paper. We added 20% of the burst duration to both ends of the burst in order to ensure that we caught the entire burst. A selection of six bursts is shown in Figure 1, to emphasize the diversity of burst morphologies we encounter.

4. Results

We computed light curves and periodograms for all 27 bursts. In each case, we produced a light curve by binning the TTE data to a time resolution of $1/2\nu_{\text{Nyquist}} = 1.22 \times 10^{-4}$ s, corresponding to a Nyquist frequency of $\nu_{\text{Nyquist}} = 4096$ Hz. We




Figure 1: Light curves of six example bursts from the magnetar SGR J0501 + 4516 recorded by Fermi/GBM. We combined data from all NaI detectors with source angles smaller than 50 degrees to the source. The time resolution corresponds to 0.005 seconds. Note the strong component of aperiodic variability after the main burst in 080823478 and the differences in peak count rate by almost one order of magnitude between the upper three bursts and the lower three.




Figure 2: Fermi/GBM observation of burst 080823847a from SGR J0501 + 4516. Left: light curve with a time resolution of 0.002 seconds. Structure in the burst profile is clearly visible. Right: unbinned periodogram (blue) and two examples of binned (magenta: 16 Hz binning; orange: 65 Hz binning) periodograms for this burst. There is a feature in the periodogram around 30 Hz (leftmost arrow), which is by itself not significant. However, significant features (arrows 2-5) are all at integer multiples of this frequency (within the uncertainty imposed by the frequency resolution), indicating the presence of harmonics at 150 Hz, 300 Hz, 900 Hz and 2100 Hz.

chose the time resolution based on the Nyquist frequency of interest: we do not expect any signals above 4000 Hz from neutron star seismic oscillations [McDermott et al. 1988]. We search both the unbinned periodogram as well as the same periodogram binned to integer multiples (3, 5, 7, 10, 15, 20, 30, 70, 100, 200, 300, 500 and 700) of the frequency resolution of that burst.

4.1. Search for QPOs

None of the 27 bursts shows periodicities of any significance in any of the unbinned periodograms. The highest data/model outlier significance is seen in burst *bn080823847a* (see Figure 2 for a light curve and periodogram), with a posterior p-value $p(T_R) = 0.11 \pm 0.01$, at frequency $\nu_{\max} = 4057$ Hz, well below the power required to reach the detection threshold

corresponding to a posterior p-value of 5%. However, the same burst shows significant signals from ~ 300 Hz to ~ 2100 Hz in the several binned periodograms (binning frequencies between 16 Hz and 1583 Hz). The most significant signals are for binning frequencies 95 Hz and 158 Hz at a frequency of ~ 2100 Hz with a posterior p-value $< 2.0 \times 10^{-5}$. Note that p-values quoted there are corrected for the number of frequencies searched, but neither for the number of bursts searched nor the number of binned spectra searched for each burst. While the former is straightforward (a simple multiplication factor of 27 for the number of bursts searched), the latter is more complicated, owing to the fact that searching different binnings for a single periodogram does not result in independent trials. The most conservative assumption is to consider them independent, including another multiplication factor equal to the number of binnings searched (here: 9). This would rule out all but the two signals with frequency bins of 95 and 158 Hz, which remain significant even after a correction for the number of trials. A possible interpretation of the observed signals follows in Section 5.

4.2. Broadband Variability

The broadband variability observed in the bursts is not just a nuisance when searching for (quasi-) periodicities, but is of interest in its own right: it shows that something is varying in the source, although not periodically. Almost all bursts in the sample are well-modeled by a simple power law. The distribution of indices ranges from 1.7 to 4.3 and peaks around 2.5, which is higher than commonly seen for example in Gamma-ray bursts (see e.g. Beloborodov et al. 2000 and Guidorzi et al. 2012). Only two bursts required the more complex model. While these were not the longest bursts, they had the highest fluence (except for the excluded saturated bursts), indicating a potential correlation between power spectral shape and burst fluence. A reliable characterisation of the broadband properties will be deferred to a future paper involving a larger sample of bursts.

5. Discussion

Magnetar bursts are a potential window into the interior of neutron stars, via the oscillations measured in magnetar giant flares. Finding analogous signals in the wealth of short SGR bursts, however, poses something of a challenge. We have shown that timing analysis of astrophysical transients is a non-trivial problem. Standard Fourier techniques are insufficient for phenomena with diverse light curve morphologies, especially when involving correlated noise processes.

Monte Carlo simulations of light curves fail to be predictive when there is no precise knowledge of the underlying burst light curve: there is a degeneracy between the overall, aperiodic burst shape, a potential red noise component, and the very thing we would like to measure: a QPO.

In the absence of better knowledge about the emission processes in magnetar bursts, we advocate a conservative Bayesian method that models the burst light curve as a pure red noise process, at the cost that weak signals are likely missed. This is the greatest weakness of our approach. Even strong signals may be undetectable at low frequencies, where burst envelope and red noise dominate. This limitation is in part not only due to restrictions of our method, but also to the short lengths of the SGR bursts, where at these frequencies only one or two cycles may be seen in the light curve. However, at frequencies close to and above 100 Hz, sensitivities approach the white noise limit, which is strongly dependent on the number of photons from a particular burst. Thus, for a bright burst with good count statistics, sensitivities are quite constraining, down to less than 10 %. This is comparable to what was observed in giant flares: for example, a QPO at 93 Hz, as seen in the 2004 flare, at roughly 10 % rms amplitude [Israel et al. 2005, Watts and Strohmayer 2006], should be detectable in at least the brightest bursts of our sample.

However, QPOs in SGR bursts may be less strong than in the giant flares, owing to the lower energy injected in SGR bursts, and hence more likely to be misclassified as non-detections, if their fractional rms amplitudes fall below 5 %.

The burst *080823847a* presents an interesting case that illustrates the limits of a pure signal-processing approach to the timing analysis shown here. The nature of the significant signals is at present unclear. They are possibly harmonics of a lower-frequency signal around 30 Hz, corresponding to a timescale of $\tau = 1/\nu = 33$ ms. This timescale roughly corresponds to the two sharp peaks seen in the burst light curve in Figure 2 (left side). Whether we consider this to be a QPO atop a burst envelope or not cannot be answered from Fourier analysis alone; it becomes a matter of interpretation and prior knowledge. At present, without any knowledge about emission processes and the kind of light curve they produce, it is impossible to distinguish whether the two sharp peaks are indeed a heavily damped QPO, or simply a chance occurrence of red noise features, thus we choose the conservative approach and interpret the observed feature as part of a noise process.

We wish to note that methods presented here, while developed with SGR bursts in mind, are by no means limited to magnetars. They are applicable in fairly general circumstances, for any light curve that is phenomenologically similar to what we observe from magnetars: highly variable, transient events with

complex light curves. This includes, for example, other known transients such as gamma-ray bursts (GRBs), tidal disruption events and supernova light curves.

Acknowledgments

D.H. and A.L.W acknowledge support from a Netherlands Organization for Scientific Research (NWO) Vidi Fellowship (PI A. Watts), and would like to thank Jason Farquhar for useful discussions. C.K. and was partially supported by NASA grant NNN07ZDA001-GLAST. E.G. acknowledges support from the Scientific and Technological Research Council of Turkey (TÜBİTAK) through grant 109T755.

References

- A. M. Beloborodov, B. E. Stern, and R. Svensson, *Astrophysical Journal* **535**, 158 (2000).
- R. C. Duncan, *Astrophysical Journal Letters* **498**, L45 (1998), arXiv:astro-ph/9803060.
- A. M. El-Mezeini and A. I. Ibrahim, *Astrophysical Journal Letters* **721**, L121 (2010), 1008.3870.
- D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *ArXiv e-prints* p. 1202.3665 (2012).
- D. W. Fox, W. H. G. Lewin, R. E. Rutledge, E. H. Morgan, R. Guerriero, L. Bildsten, M. van der Klis, J. van Paradijs, C. B. Moore, T. Dotani, et al., *Monthly Notices of the Royal Astronomical Society* **321**, 776 (2001), arXiv:astro-ph/0009224.
- E. Göğüş, P. M. Woods, C. Kouveliotou, J. van Paradijs, M. S. Briggs, R. C. Duncan, and C. Thompson, *Astrophysical Journal Letters* **526**, L93 (1999), arXiv:astro-ph/9910062.
- E. Göğüş, P. M. Woods, C. Kouveliotou, J. van Paradijs, M. S. Briggs, R. C. Duncan, and C. Thompson, *Astrophysical Journal Letters* **532**, L121 (2000), arXiv:astro-ph/0002181.
- C. Guidorzi, *Monthly Notices of the Royal Astronomical Society* **415**, 3561 (2011), 1104.5308.
- C. Guidorzi, R. Margutti, L. Amati, S. Campana, M. Orlandini, P. Romano, M. Stamatikos, and G. Tagliaferri, *Monthly Notices of the Royal Astronomical Society* **422**, 1785 (2012), 1202.3443.
- D. Huppenkothen, A. L. Watts, P. Uttley, A. J. van der Horst, M. van der Klis, C. Kouveliotou, E. Gogus, J. Granot, S. Vaughan, and M. H. Finger, *Astrophysical Journal*, in press (2013), 1212.1011.
- K. Hurley, T. Cline, E. Mazets, S. Barthelmy, P. Butterworth, F. Marshall, D. Palmer, R. Aptekar, S. Golenetskii, V. Il'Inskii, et al., *Nature* **397**, 41 (1999), arXiv:astro-ph/9811443.
- G. L. Israel, T. Belloni, L. Stella, Y. Rephaeli, D. E. Gruber, P. Casella, S. Dall'Osso, N. Rea, M. Persic, and R. E. Rothschild, *Astrophysical Journal Letters* **628**, L53 (2005), arXiv:astro-ph/0505255.
- Y. Kaneko, E. Göğüş, C. Kouveliotou, J. Granot, E. Ramirez-Ruiz, A. J. van der Horst, A. L. Watts, M. H. Finger, N. Gehrels, A. Pe'er, et al., *Astrophysical Journal* **710**, 1335 (2010), 0911.4636.
- C. Kouveliotou, S. Dieters, T. Strohmayer, J. van Paradijs, G. J. Fishman, C. A. Meegan, K. Hurley, J. Kommers, I. Smith, D. Frail, et al., *Astrophysical Journal Nature* **393**, 235 (1998).
- C. Kouveliotou, T. Strohmayer, K. Hurley, J. van Paradijs, M. H. Finger, S. Dieters, P. Woods, C. Thompson, and R. C. Duncan, *Astrophysical Journal Letters* **510**, L115 (1999), arXiv:astro-ph/9809140.
- J. M. Lattimer and M. Prakash, *Physics Reports* **442**, 109 (2007), arXiv:astro-ph/0612440.
- L. Lin, C. Kouveliotou, M. G. Baring, A. J. van der Horst, S. Guiriec, P. M. Woods, E. Göğüş, Y. Kaneko, J. Scargle, J. Granot, et al., *Astrophysical Journal* **739**, 87 (2011).
- P. N. McDermott, H. M. van Horn, and C. J. Hansen, *Astrophysical Journal* **325**, 725 (1988).
- C. Meegan, G. Lichti, P. N. Bhat, E. Bissaldi, M. S. Briggs, V. Connaughton, R. Diehl, G. Fishman, J. Greiner, A. S. Hoover, et al., *Astrophysical Journal* **702**, 791 (2009), 0908.0450.
- S. Mereghetti, *The Astronomy and Astrophysics Review* **15**, 225 (2008), 0804.0250.
- D. M. Palmer, S. Barthelmy, N. Gehrels, R. M. Kippen, T. Cayton, C. Kouveliotou, D. Eichler, R. A. M. J. Wijers, P. M. Woods, J. Granot, et al., *Nature* **434**, 1107 (2005), arXiv:astro-ph/0503030.
- T. E. Strohmayer and A. L. Watts, *Astrophysical Journal Letters* **632**, L111 (2005), arXiv:astro-ph/0508206.
- T. E. Strohmayer and A. L. Watts, *Astrophysical Journal* **653**, 593 (2006), arXiv:astro-ph/0608463.
- C. Thompson and R. C. Duncan, *Monthly Notices of the Royal Astronomical Society* **275**, 255 (1995).
- S. Vaughan, *Monthly Notices of the Royal Astronomical Society* **402**, 307 (2010), 0910.2706.
- A. L. Watts and T. E. Strohmayer, *Astrophysical Journal Letters* **637**, L117 (2006), arXiv:astro-ph/0512630.
- A. L. Watts, in *Neutron Star Crust*, edited by C. A. Bertulani and J. Piekarewicz (Nova Science Pub Inc, 2012), pp. 265–280 arXiv:astro-ph/1111.0514.
- P. M. Woods, C. Kouveliotou, E. Göğüş, M. H. Finger, J. Swank, D. A. Smith, K. Hurley, and C. Thompson, *Astrophysical Journal* **552**, 748 (2001), arXiv:astro-ph/0101045.
- P. M. Woods and C. Thompson, in *Compact stellar X-ray sources*, edited by W. H. G. Lewin and M. van der Klis (Cambridge University Press, 2006), pp. 547–586.